

SYSTEMATIC ERRORS IN LOW-LATENCY GRAVITATIONAL WAVE PARAMETER ESTIMATION IMPACT ELECTROMAGNETIC FOLLOW-UP OBSERVATIONS

TYSON B. LITTENBERG

Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Department of Physics and Astronomy,
 Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

BEN FARR

Enrico Fermi Institute, Department of Physics, and Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637,
 USA

SCOTT COUGHLIN

Cardiff University, Cardiff CF24 3AA, UK

VICKY KALOGERA

Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Department of Physics and Astronomy,
 Northwestern University, 2145 Sheridan Road, Evanston, IL 60208

(Dated: March 2, 2016)
 Draft version March 2, 2016

ABSTRACT

Among the most eagerly anticipated opportunities made possible by Advanced LIGO/Virgo are multimessenger observations of compact mergers. Optical counterparts may be short-lived so rapid characterization of gravitational wave (GW) events is paramount for discovering electromagnetic signatures. One way to meet the demand for rapid GW parameter estimation is to trade off accuracy for speed, using waveform models with simplified treatment of the compact objects' spin. We report on the systematic errors in GW parameter estimation suffered when using different spin approximations to recover generic signals. Component mass measurements can be biased by $> 5\sigma$ using simple-precession waveforms and in excess of 20σ when non-spinning templates are employed. This suggests that electromagnetic observing campaigns should not take a strict approach to selecting which LIGO/Virgo candidates warrant follow-up observations based on low-latency mass estimates. For sky localization, we find searched areas are up to a factor of ~ 2 larger for non-spinning analyses, and are systematically larger for any of the simplified waveforms considered in our analysis. Distance biases for the non-precessing waveforms can be in excess of 100% and are largest when the spin angular momenta are in the orbital plane of the binary. We confirm that spin-aligned waveforms should be used for low-latency parameter estimation at the minimum. Including simple precession, though more computationally costly, mitigates biases except for signals with extreme precession effects. Our results shine a spotlight on the critical need for development of computationally inexpensive precessing waveforms and/or massively parallel algorithms for parameter estimation.

1. INTRODUCTION

Ground-based gravitational wave (GW) interferometers are most sensitive at frequencies between 10 Hz and 2 kHz. Binary systems comprising of compact stellar remnants such as black holes (BHs) and neutron stars (NSs) will experience the late stages of orbit inspiral and will merge at these frequencies, making them the prime target for the Advanced LIGO (Abadie *et al.* 2014) and Advanced Virgo (Acernese *et al.* 2015) network of detectors. The scientific payoff from observing a population of compact binaries in GWs is hard to overstate. Mergers of binary neutron star (BNS) and/or neutron star-black hole (NSBH) systems are the leading candidates for the progenitors of short-hard gamma-ray bursts (sGRBs). Comparing rates of GW detections and sGRBs will constrain the beaming angle for the gamma-ray jet while GW signatures from such systems coincident with GRBs will confirm the binary merger scenario (Chen and Holz 2013). Details of the inspi-

ral and merger signals will probe regimes of strong-field, dynamical gravity inaccessible to any other observational or experimental techniques, testing the general relativistic field equations (Cornish *et al.* 2011; Li *et al.* 2012; Sampson *et al.* 2013; Berti *et al.* 2015) and examining the equation of state of matter at supranuclear density (Flanagan and Hinderer 2008; Read *et al.* 2009; Lackey *et al.* 2014; Del Pozzo *et al.* 2013; Wade *et al.* 2014). The merger rates and physical properties of compact binaries are sensitive to physical processes of compact object formation and binary evolution; their measurements will provide strong constraints on theoretical models for population synthesis (Belczynski *et al.* 2010, 2012; Mandel *et al.* 2015). Last but not least, well localized binaries will serve as cosmological “standard sirens,” further constraining the scale and expansion rate of the local universe (Schutz 1986; Holz and Hughes 2005; Nissanke *et al.* 2010).

Clearly the advent of GW astronomy and astrophysics

is highly anticipated in its own right. Furthermore, GW discoveries with identified electromagnetic (EM) counterparts would allow us to reach the full potential of multimessenger astronomy. GW observatories are all-sky monitors and are well approximated by a quadrupole antenna with maximum sensitivity directly above and below the plane of the detector. Because the antenna pattern varies slowly with angular distance away from the detector plane zenith, single detectors have poor sky-localization capabilities. Differences in the arrival time, phase, and polarization of the GW signal in different detectors separated by large baselines enable more precise sky-location measurements because a relatively small region of the sky could produce a particular coherent signal.

Because the GW measurement will not be able to pinpoint the location of a source alone, great effort is needed to conduct electromagnetic observing campaigns. Because any optical counterpart may be short-lived, and a typical EM telescope will have to use several pointings to cover the LIGO/Virgo error region on the sky, time is of the essence. Unfortunately, obtaining the most precise and accurate estimates for a GW event’s physical parameters is a time-consuming computational challenge. LIGO/Virgo parameter estimation is performed using stochastic Bayesian sampling algorithms such as Markov Chain Monte Carlo (MCMC) (Gelman 1997), Nested Sampling (Skilling 2004), or MultiNest (Feroz and Hobson 2007). Circularized binary inspiral signals are characterized by 15 parameters. The 15-dimensional posterior distribution function is highly structured, often exhibiting multiple modes and strong correlations between parameters. The sampling algorithms perform a fundamentally serial random exploration of parameter space, at each point evaluating how well a trial signal matches the data. Each step of the analysis requires the computation of a model GW waveform, or “template,” which must be repeated $\mathcal{O}(10^7)$ times to sufficiently characterize the 15-dimensional parameter space. The standard sampling algorithms take $\mathcal{O}(\text{days})$ of modern high-performance computing time to converge—latencies which can be longer in duration than the lifetime of proposed EM signatures such as kilonova (Metzger *et al.* 2010).

The two philosophical approaches to reducing the real-time latency for parameter estimation are to use computationally inexpensive approximations to the template models or a simplification of the likelihood function. Several groups have explored novel, massively parallel, parameter estimation strategies to further reduce the time required for low-latency parameter estimation (Pankow *et al.* 2015; Haster *et al.* 2015; Farr 2015). While the massively parallel parameter estimation methods have shown promise in providing prompt, robust, signal characterization none have reached sufficient maturity to be deployed in the real-time analysis for the early Advanced LIGO observing runs.

GW parameters are often divided into *extrinsic* parameters (orientation, sky location, distance), which depend on the observer, and *intrinsic* parameters (masses and spins of the compact objects), which depend on the binary system. Generally speaking, intrinsic parameters influence the phase evolution of the binary waveform while extrinsic parameters are responsible for the rela-

tive amplitude, polarization, and time of arrival of the signal at each detector. To meet the low-latency demands placed on GW parameter estimation for prompt EM follow-up, a hierarchical strategy has been devised that takes advantage of the separability between intrinsic and extrinsic parameters to rapidly report parameters for a candidate GW event (see Farr *et al.* (2014) for discussion on separate intrinsic and extrinsic analyses). In this paradigm computation time is saved by reducing the number of intrinsic degrees of freedom in the GW model by making simplifying assumptions about the spin of the compact objects, either assuming particular alignments for the spin angular momenta or fixing these terms to be zero. The scheme first relies on the relative signal-to-noise ratio (S/N) and phase in each detector with latency of minutes (Singer and Price 2015), a parameter estimation analysis where the templates are restricted by assuming either non-spinning or aligned spin binaries (substantially decreasing run time to a few hours) (Singer *et al.* 2014; Berry *et al.* 2015); finally the full parameter estimation analysis that uses generically spinning waveforms (Farr *et al.* 2015).

The separability of intrinsic and extrinsic parameters is a fair approximation when dealing with binary NS systems where spins are expected to be small, possibly less than 10% of the maximum value allowed by general relativity. This separability breaks down when the binary contains one or two BHs, as their spins can be near maximal (see Fragos and McClintock (2015) and references therein). Through relativistic couplings between the spin-spin and spin-orbit angular momenta, rapidly spinning misaligned binaries undergo dramatic precession of the orbital plane, thereby coupling the intrinsic and extrinsic parameters. As a result, restricting the spin degrees of freedom for template waveforms can bias the parameters we wish to measure rapidly, namely sky location of the GW source and the distance to a binary signal, as well as to whether the source involves a NS or not (the latter has implications for the potential existence of an EM counterpart).

Miller *et al.* (2015) systematically compared the results of parameter estimation from using different waveforms, varying different spin parameterizations as well as post-Newtonian families, finding that biases in intrinsic parameter estimation (i.e. masses and spins) are comparable to the systematic differences between post-Newtonian families; they did not address issues related to extrinsic parameters of importance, such as sky location and distance.

In this paper, using a consistent post-Newtonian treatment for all waveforms, we put the biases due to reducing the spin degrees of freedom into astrophysical context. We address the following key question: do the spin-constrained, fast, frequency-domain waveforms planned for low-latency analysis introduce important biases in parameter estimation, specifically relevant to EM follow-up work (which is the reason for low-latency analysis)? We find that systematic errors incurred by using non-spinning or spin-aligned templates will be significant, with respect to both the statistical errors and the astrophysical interpretation of the results. For component masses that may be used to determine which candidate events should be followed up with telescopes, the biases can be in excess of several M_{\odot} , opening the possibility

of misidentifying a NS as a BH or vice versa. Electromagnetic searches will have to cover an area on the sky larger by a factor of $\lesssim 2$ to have the same probability of pointing at the true source location compared to the sky localization of the fully spinning/precessing analysis. While the biases systematically decrease as we use increasingly complete waveforms, we find no substitute, fast, frequency-domain template waveform that is immune to introducing substantial biases in some parts of parameter space.

In the near term our results should be used to guide follow-up observing strategies. For the future we hope that our finding helps stimulate further development of fast, accurate waveforms that contain all of the necessary degrees of freedom for accurate parameter estimation.

2. SETUP

We simulate 1000 LIGO/Virgo signals randomly selected from a uniform distribution in component mass (m_1, m_2) drawn from $m \in [1, 30] M_\odot$ with an additional constraint on the total mass of the binary $m_1 + m_2 \leq 30 M_\odot$; six parameters describing the two spin vectors $\mathbf{S}_1, \mathbf{S}_2$ with magnitudes drawn from $U[0, 1]$ and isotropic orientation with respect to the orbital angular momentum \mathbf{L} ; sky location (α, δ) distributed uniformly on the celestial sphere; and three angles drawn uniformly over a sphere that describe the orientation of the binary $(\theta_{J,N}, \psi, \phi_0)$. The $\theta_{J,N}$ parameter is the angle between the wave propagation direction \hat{N} (opposite the observer’s line of sight \mathbf{k}) and the total angular momentum \mathbf{J} . In the limit of non-spinning binaries this is degenerate with the inclination angle $\iota = \cos^{-1}(\hat{k} \cdot \hat{L})$ which is more traditionally used in the GW literature. We prefer $\theta_{J,N}$ because it has less variation during the orbital evolution, while the direction of \mathbf{L} is time-varying for spin configurations that induce significant precession (Farr *et al.* 2014). The luminosity distance, D_L , of each source is drawn uniformly in volume with an additional detectability requirement that the binary has a signal-to-noise ratio $S/N > 5$ in at least two detectors. This S/N requirement produces a strong selection effect due to the quadrupolar emission of gravitational radiation. While GW emission from a compact binary is omnidirectional, it is strongest along the direction of the orbital angular momentum vector \mathbf{L} of the binary. A system that is “face on,” i.e. the line of sight to the observer \mathbf{k} and \mathbf{L} , is nearly parallel ($\hat{k} \cdot \hat{L} \sim 1$), is detectable out to a larger distance than the same system in an “edge on” ($\hat{k} \cdot \hat{L} \sim 0$) configuration. Assuming that compact binaries are distributed uniformly in volume, far more face-on binaries will be detected than edge-on binaries (see Fig. 2 and related discussion in Littenberg *et al.* (2015)).

In our simulations, we generate the source GW signals using the SpinTaylorT2 waveforms, which are time-domain inspiral-only waveforms that allow for generic spin magnitude and orientation, and hence for full precession. The waveforms do not include any treatment of the merger or ringdown signal, and assume quasi-circular orbits and negligible finite-size effects of the components. For each binary we compute the response of Advanced LIGO/Virgo at design sensitivity Aasi *et al.* (2013) with a low-frequency cutoff at 20 Hz. Each simulated signal is processed by the LIGO/Virgo data analysis pipeline

for parameter estimation, LALInference. LALInference is a code library built to provide robust Bayesian analysis of GW signals and is part of the LIGO Analysis Library (LAL). For this study we use the LALINFERENCE_MCMC application in LALInference, a parallel tempered MCMC algorithm. A full description of the LALInference package, and the custom features found in `lalinference_mcmc` for analysis of compact binaries, can be found in Veitch *et al.* (2015).

The LALInference analysis is repeated using four template waveforms: SpinTaylorT2, which provides bias-free parameter estimation of our simulated signals (“full precession”), and three other faster, frequency-domain waveforms, all inspiral-only, for consistency with the source signals: SpinTaylorF2, which uses a simplified treatment of the precession that conserves the opening angle $\beta = \cos^{-1}(\hat{J} \cdot \hat{L})$ (“simple precession”); TaylorF2Aligned, which fits for spin magnitudes assuming $\beta = 0$ (“aligned spin”); and zero-spin TaylorF2 waveforms (“non-spinning”). Figure 1 shows the computational cost per likelihood evaluation between the different waveforms employed in this work. The approximate cost for each waveform model is determined by using the `lalinference_mcmc` software to draw 1000 independent samples for the prior distribution. The results are normalized by the computation time for the non-spinning TaylorF2 waveforms because they are the fastest available and the baseline model for rapid parameter estimation. Including the two spin magnitudes in a constrained, fully aligned configuration (TaylorF2aligned) has a negligible effect on the computational cost of the model. Allowing for simple precession (SpinTaylorF2) comes with a factor of ~ 2 increase in run time, while the time-domain, fully precessing waveform (SpinTaylorT2) is over an order of magnitude more computationally taxing. This simple benchmarking test does not take into account the increased difficulty of sampling higher dimension posteriors, which can increase the convergence time for the MCMC analysis, but serves as a useful reference for the tradeoff between waveform accuracy and speed.

The output of `lalinference_mcmc` is a series of independent samples drawn from the posterior distribution function for the template parameters that are used to draw inferences about the astrophysical nature of the GW source.

3. RESULTS: INTRINSIC PARAMETERS

We begin by focusing on intrinsic parameter biases. Because the spin and mass parameters govern the phase evolution of the binary, artificially constraining spin degrees of freedom for generic binaries will have a dramatic effect on mass measurement. Figure 2 summarizes the consequences of restricting spin degrees of freedom for mass determination. Results shown in this figure come from marginalized posterior distributions for the smaller mass in the binary, m_2 . We use m_2 as a proxy for whether a GW signal is a good candidate for electromagnetic observations. While BNS and NSBH systems are considered possible progenitors of short-duration GRBs, there are no strong theoretical candidates from electromagnetic signals from merging stellar-mass BHs. A resource-limited electromagnetic follow-up campaign may consider using the inferred m_2 to determine which candidate




FIG. 1.— Relative speed for different waveforms tested in this work estimated by drawing 1000 independent samples from the `lalinference_mcmc` prior. All times are measured relative to the TaylorF2 (non-spinning) waveforms. Spin-aligned (TaylorF2Aligned) waveforms add negligible cost. Simple precession (SpinTaylorF2) is a factor of ~ 2 more costly, while the full precession waveform (SpinTaylorT2) is slower by an order of magnitude. These results do not take into account the added inefficiency of sampling the higher dimension posteriors.

LIGO/Virgo detections are worth telescope time, selecting only those systems that have an m_2 measurement consistent with NS masses.

The left panel is a histogram of the bias on m_2 . The bias is measured as the difference Δm_2 between the mean of the m_2 posterior and the true value, normalized by the posterior’s standard deviation σ . Using this definition, the bias is a measure of how many standard deviations (σ) away from the true value the posterior distribution is peaked. In this figure and throughout this paper, the gray, blue, green, and red curves correspond to the precessing, simple precession, aligned spin, and non-spinning waveforms, respectively. In Fig. 2 we find that the possible biases can be in excess of 5σ even when using the simple-precession waveforms, and reach beyond 20σ for the non-spinning waveforms often used for low-latency parameter estimation. The right panel is the distribution of standard deviations for the same simulated signals, which are strongly peaked around $1 M_\odot$ with a long tail down to $\sim 10^{-2} M_\odot$ for the non-spinning waveforms. Not only is a 10σ or 20σ bias statistically alarming, but because of the typical standard deviations for this parameter, these biases can have astrophysical repercussions as well (Littenberg *et al.* 2015).

Figure 3 shows 11 randomly selected binaries (open circles) from our simulated population with a scatter plot of the m_1 – m_2 posterior samples for each source. The heavy vertical and horizontal lines denote the putative “mass gap” between the theorized highest mass NSs ($\sim 3 M_\odot$) and the observed lowest mass BHs ($\sim 5 M_\odot$) (Farr *et al.* 2011; Özel *et al.* 2010; Kreidberg *et al.* 2012). Each panel shows the recovered masses for the different waveforms tested superimposed with the bias-free full precession posteriors, which are shown in gray in each panel, and alone in the top left. Here we find that the non-spinning waveforms (red, bottom right) are highly prone to misclassifying the GW source. The true

value (open circle) is typically well outside of the region in m_1 – m_2 space covered by the posterior and often confined in the wrong region of parameter space entirely (e.g. in the NSBH region when the true signal was a binary BH of comparable mass). When including some freedom in the spin parameters, whether that be the aligned (green, bottom left) or simple precession case (blue, top right), we find more consistent overlap with the bias-free posteriors. However, neither the aligned spin nor the simple-precession waveforms can generically be substituted for the full precession waveforms without risking some substantial biases in the determination of m_1 and m_2 . For example, both the aligned spin and simple-precession waveforms misidentified a source with a low-mass secondary ($m_1 \sim 19M_\odot$ and $m_2 \sim 4.5M_\odot$) as a binary BH system of comparable mass. We conclude that low-latency mass estimates are not generically reliable, and that electromagnetic observing campaigns should not take a strict approach to selecting which LIGO/Virgo candidates warrant follow-up observations based on early mass estimates.

4. RESULTS: EXTRINSIC PARAMETERS

The most pertinent measurements needed for low-latency GW searches are the sky location and distance to the source. The sky location is of obvious importance for follow-up electromagnetic observations, and the distance can be used to either select nearby candidate GW sources for observation or to convolve the GW measurement with galaxy catalogs (Nissanke *et al.* 2013; Hanna *et al.* 2014; Singer *et al.* 2014; Fan *et al.* 2014).

For sky localization, our figure of merit is the “searched area” for different waveforms. The searched area is a measure of the number of square degrees contained within the credible interval of the posterior that intersects the true signal parameters. If an electromagnetic follow-up strategy is to sequentially point a telescope at fields on the sky ranked by their posterior weight, the searched area gives a sense of how large a region needs to be searched to have a certain probability of the true source location being included among the observations.

Figure 4 shows the cumulative distribution function of searched area for the precessing, simple precession, aligned spin, and non-spinning waveforms. We see that searched areas are systematically shifted to larger values as the spin degrees of freedom are further restricted, meaning that sky-location biases incurred by the simplified waveforms translate to more telescope pointings to ensure the same probability of imaging the true source. For most signals in our simulated population, the differences in searched area between the fully precessing waveforms and the non-spinning waveforms are around a factor of two. Using spin-aligned waveforms yields searched areas that are typically $\sim 30\%$ to $\sim 50\%$ larger than the full-precession results. Simple-precession searched areas are, at most, $\sim 30\%$ larger but typically are within $\sim 20\%$ of the full-precession areas.

The distance to the source is a more difficult parameter to measure. As noted throughout the GW literature, first in Cutler and Flanagan (1994) and further investigated in Nissanke *et al.* (2010) and Rodriguez *et al.* (2014), there is a strong degeneracy between the orientation of the binary and the distance which causes the waveforms of near by, edge-on systems to be nearly in-




FIG. 2.— Left: bias in component mass recovery for full precession (gray, solid), simple precession (blue, dashed), aligned spin (green, dotted), and non-spinning (red, dotted-dashed) templates measured in standard deviations of the full-precession posterior. We find large biases when using simplified treatments of the spin parameters. Right: same color scheme as the left panel, now showing the distribution of the standard deviation for the component mass posteriors, which are typically $\sim 1 M_{\odot}$. The biases we find are astrophysically significant.

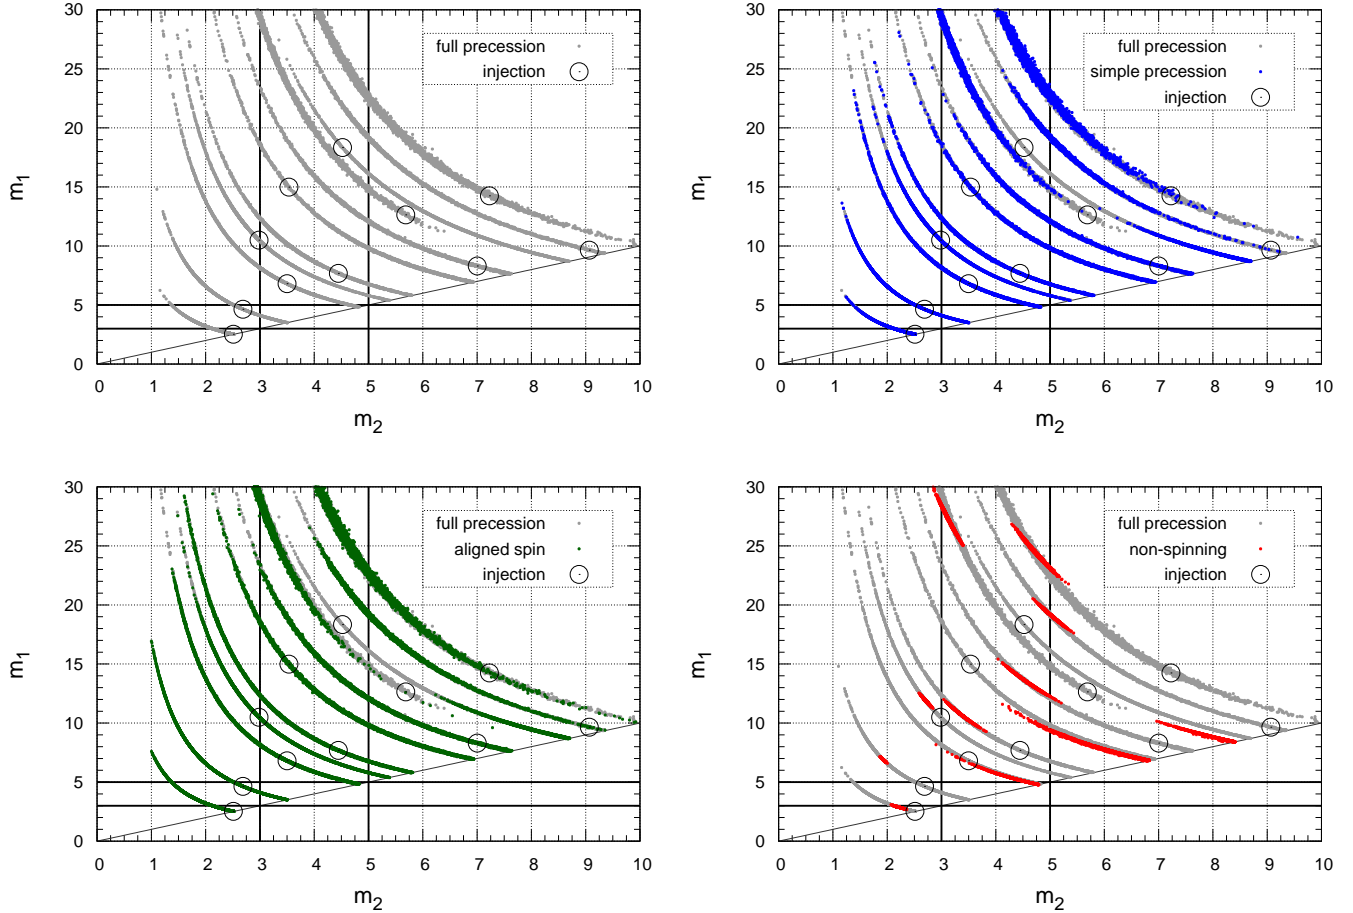


FIG. 3.— A selection of m_1 - m_2 posteriors compared to the true source values (open circles) for the four different spin configurations: full precession (gray points in each panel, plotted by themselves in the top left); simple precession (blue, top right); aligned spin (green, bottom left); and non-spinning (red, bottom right). Heavy vertical and horizontal lines mark the “mass gap” boundary between high-mass NSs and low-mass BHs. The inferred component masses from non-spinning waveforms generically miss the true source value and often incorrectly classify the binary. Including aligned spins or simple precession improves the agreement with the full-precession results but misclassification is still possible. Low-latency mass estimates should be taken with caution.




FIG. 4.— Cumulative distribution functions for the searched sky area of each injection when using the non-spinning templates (red, dotted-dashed), aligned spin (green, dotted), simple precession (blue, dashed), and the precessing templates (gray, solid). The non-spinning analysis systematically requires more searched area to find the true signal. Including more accuracy in the waveform systematically improves searched areas. None of the simplified waveforms matches the performance of the fully precessing templates.

distinguishable from distant, face-on systems at plausible S/Ns for Advanced LIGO and Virgo observations. Convolution with the assumption that binaries are uniformly distributed within the LIGO/Virgo horizon results in distance posteriors heavily influenced by the prior, which strongly favors distant, face-on systems.

To quantify the difference in distance measurement we use the bias-free fully precessing distance estimates as the baseline against which the other waveforms’ performances are measured. Figure 5 shows the fractional difference $\Delta D_L/D_L$ in distance D_L , defined as the difference between the peaks of the distance posteriors for the waveform under investigation and the full-precession waveform divided by the peak of the full-precession distance posterior. The error bars in Figure 5 are one standard deviation of the posterior normalized by the peak of the full-precession distance distribution. The result is a measure of the bias and fractional error in distance measurement from using waveforms that make simplifying approximations to the spin. From top to bottom, the rows of plots correspond to the full precession (gray), simple precession (blue), aligned spin (green), and non-spinning (red) templates used for recovery. Our use of the mode of the distance posterior as the point estimate against which we make comparisons was for convenience. Macroscopic differences between the distributions are insensitive to the particular choice of definition for bias in distance.

The left column in Figure 5 shows $\Delta D_L/D_L$ as a function of the true signal’s network S/N. When using the same waveform for recovery as we do for simulating the data (top row) we find statistical errors of 10-30%. We find that the distance biases grow and exceed the statistical uncertainty as the waveforms employed for recovery use simpler treatment of the spin. The biases in distance determination are not found to be strongly dependent on S/N. As a consequence, for accurate distance estimation spin precession will have to be considered for any

detectable signal.

The right column in Figure 5 presents the fractional error in distance $\Delta D_L/D_L$ as a function of the opening angle β between total and orbital angular momenta. Here we see that the biases become larger than the statistical error for the aligned (third row, green) and non-spinning (bottom row, red) waveforms for $\beta > 10^\circ$ and can be larger than 100% for $\beta > 30^\circ$. At $\beta = 90^\circ$ the spin angular momentum vector lies in the orbital plane of the binary, maximizing the orbital precession. The aligned-spin and non-spinning waveforms assume a constant binary orientation, making these waveforms a poor match to the true signal. Except for rare instances the simple-precession waveforms (second row, blue) return qualitatively similar distance posteriors to the full-precession waveforms.

None of the results shown have assumed anything about the GW source beyond the implicit validity of the waveform models used for signal recovery. A special case for ground-based GW astrophysics is the simultaneous observation of the GW and EM signature from a compact merger. The most plausible EM counterparts to binary mergers involving at least one NS are short GRBs. It is generally assumed that the gamma-rays are emitted along a narrow, relativistic jet with an opening angle of order 10° . In a joint GRB/GW analysis, assuming the jet is parallel to the angular momentum vector, gamma ray detections place a strong constraint on the allowed values of $\theta_{J,N}$, and consequently D_L for the binary.

For a sense of how the inferred distance depends on different spin approximations, we select NSBH events from our population of simulated signals that had $\theta_{J,N} < 20^\circ$ (or $> 160^\circ$) as a set of potential short GRBs. If coincident GRB detections were made we could adopt a strict prior on the orientation of the binary, which would help constrain the distance to the source. Figure 6 shows 90% credible intervals for the distance when we adopt the “GRB prior” for $\theta_{J,N}$ as a function of the true signals’ S/N. The black circles denote the true distance. Once more we find significant bias in the distance estimate when non-spinning (red, dotted-dashed) waveforms are used for the analysis. The aligned spin (green, dotted) analysis excludes the true distance for three of the 10 simulation. For this small sample the simple precession (blue, dashed) and full precession (gray, solid) waveforms return qualitatively similar distance posteriors, and in only one event do they differ significantly.

5. DISCUSSION

Among the most eagerly anticipated discoveries made possible by Advanced LIGO/Virgo are joint GW and EM observations of compact mergers. Optical counterparts of compact mergers may be short-lived so a rapid response of electromagnetic follow-up observations to GW triggers is paramount for discovering the EM signature of GW detections. One way to meet the demand for rapid GW parameter estimation is to trade off accuracy for speed, using computationally inexpensive waveform models that make simplifying assumptions about the spin of the constituent compact objects.

In this paper we studied the biases in parameter estimation for generic GW signals when non-spinning, aligned-spin, and simple-precession templates are used for the analysis. We find non-negligible systematic

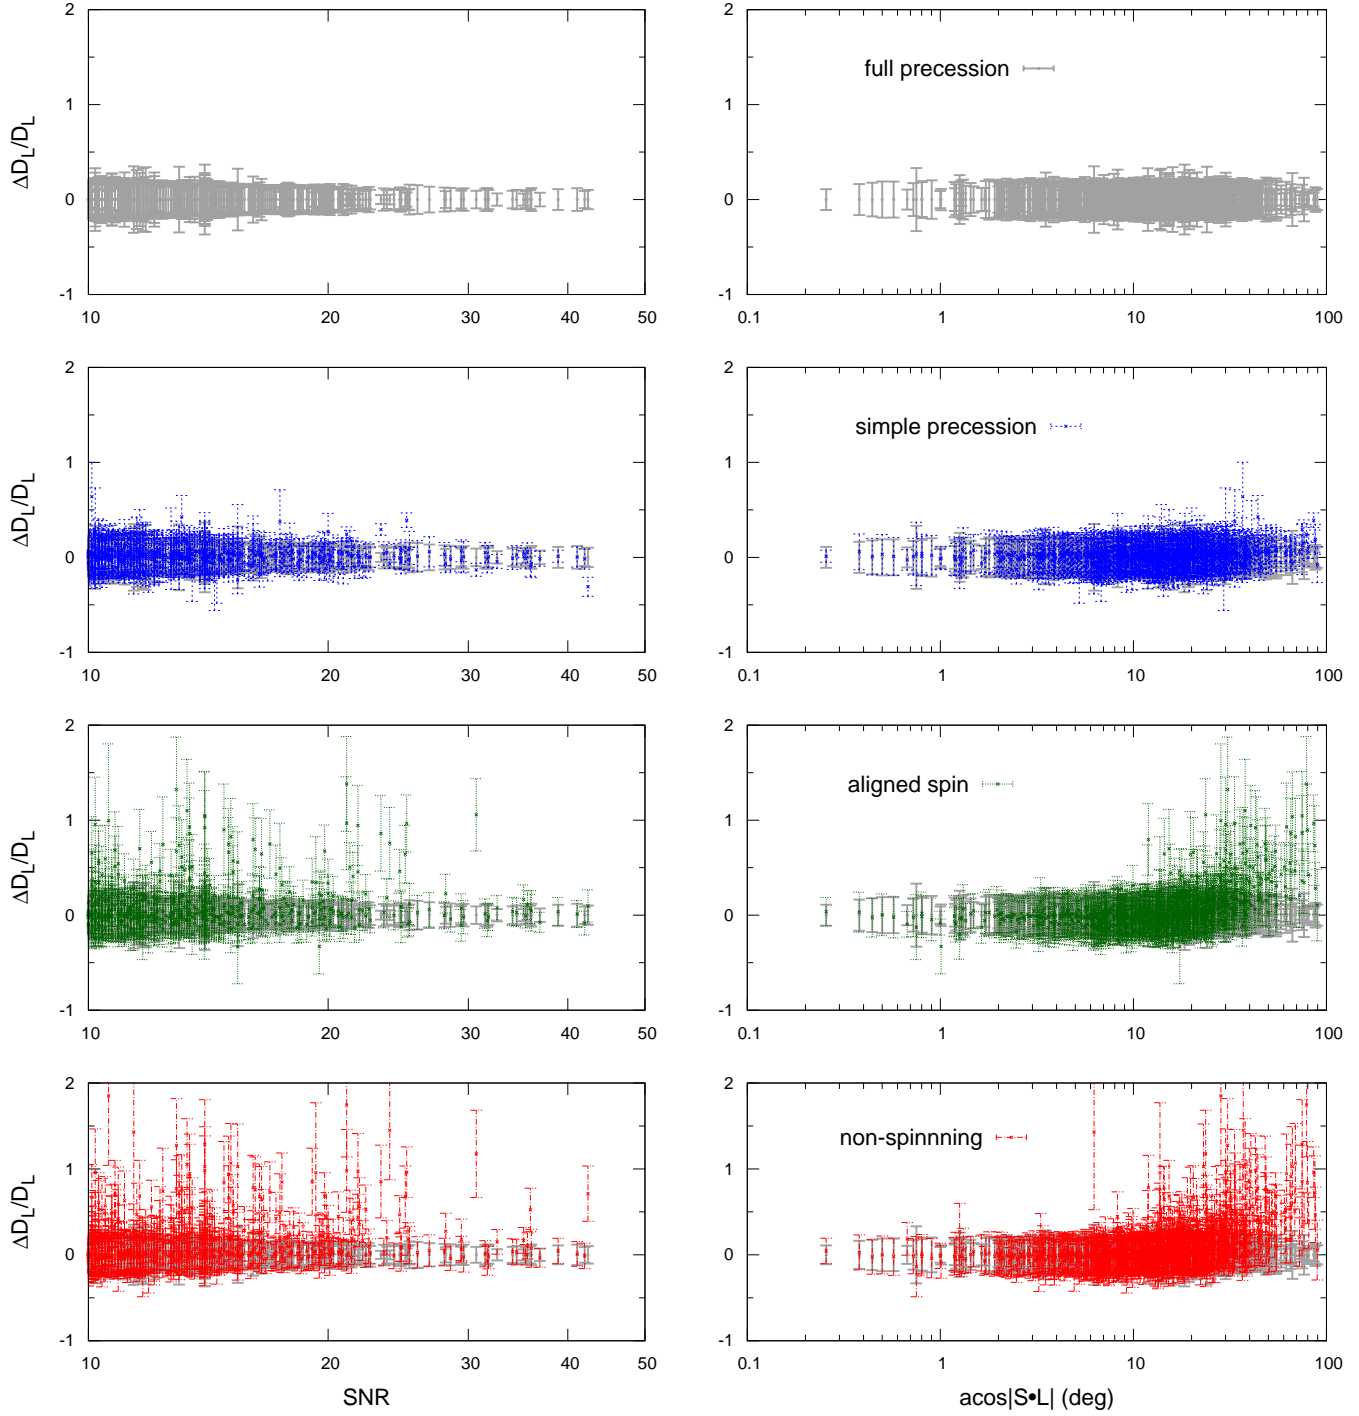


FIG. 5.— Left column: fractional differences in distance as a function of S/N. From top to bottom the results correspond to the full precession, simple precession, aligned spin, and non-spinning analyses. Systematic errors in distance are largest for waveforms without precession and the errors do not strongly scale with S/N. Right column: same ordering but now the independent variable is the opening angle β between the total and orbital angular momenta. Here we see that the bias is strongest for opening angles close to 90° , which are cases that exhibit the most dramatic precession effects.




FIG. 6.— Examples of marginalized distance posteriors for full precession (gray, solid), simple precession (blue, dashed), aligned spin (green, dotted), and non-spinning (red, dotted-dashed) distance posteriors. Results correspond to the NSBH systems in our injection set with face-on ($\theta_{J,N} < 20^\circ$) orientations. Black circles mark the true value. Points are the median values and error bars span the 90% credible interval of the posterior. For all cases the simple precession and precession recovery include the injection. The aligned-spin analysis excludes the true distance for three of the 10 simulations. The non-spinning analysis excludes the true distance in all but four of the simulations.

errors on both intrinsic (masses) and extrinsic (location and orientation) parameters. Our study builds on Miller *et al.* (2015), which investigated systematic differences between post-Newtonian waveforms, by putting the biases resulting from restricting the spin degrees of freedom into astrophysical context. We conclude that the biases in parameter estimation of spinning binaries using non-spinning templates will be large enough to impact astrophysical interpretation of the results.

Component mass measurements, which may be used to determine which candidate events should be followed up with telescopes, can be biased by $> 5\sigma$ using simple-precession waveforms and in excess of 20σ when non-spinning templates are employed. These systematic errors translate to biases several M_\odot large, potentially leading to misclassification of NSs and BHs. This suggests that electromagnetic observing campaigns should not take a strict approach to selecting which LIGO/Virgo candidates warrant follow-up observations based on low-latency mass estimates.

Without question the most pertinent parameter estimation result of GW searches is the inferred sky location of the source. We find that searched areas are up to a factor of ~ 2 larger for non-spinning analyses, and are systematically larger for any of the simplified waveforms considered in our analysis. While the searched areas systematically decrease as we use increasingly complete waveforms, we find no substitute template waveform that is able to reproduce the sky-localization performance of the fully spinning/precessing waveforms.

Distance determination can be used for selecting which candidates to observe, optimizing exposure times, and employing galaxy catalogs to help guide EM observations. Distance biases for the non-precessing waveforms can be in excess of 100% and are largest when the spin angular momenta are in the orbital plane of the binary. The simple-precession waveforms produce consistently similar distance posteriors to the full-precession waveforms and are the only reliable low-latency option if distance estimates are deemed important for electromagnetic follow-up observations. For potential GRB sources in our simulated population—NSBH injections with orbital angular momentum oriented within 20° —the simple-precession waveforms produced distance estimates consistent with the full-precession waveforms while the aligned-spin recovery excluded the true distance at the 90% level for three of the 10 injections. The non-spinning waveforms showed a substantial bias, often overestimating the distance to the source by hundreds of Mpc.

In the near term, we confirm that spin-aligned waveforms should be used for low-latency parameter estimation at the minimum. The computational cost is nearly identical to that of the non-spinning waveforms, and the inclusion of spin dramatically improves mass estimates and distance estimates for face-on systems. The aligned-spin models are still subject to significant biases in the distance when the spin angular momenta are in the orbital plane of the binary, and yield larger searched areas than the fully precessing waveforms. At the added cost of a factor of ~ 2 in computational effort, the simple-precession waveforms correct biases in the distance and make further improvements to the sky localization. However, at this time there is no substitute for the full spinning precession analysis for accurate parameter estimation in any of the parameters we studied. Our results shine a spotlight on the importance of efforts to develop computationally inexpensive precessing waveforms (Field *et al.* 2014; Lundgren and O’Shaughnessy 2014; Chatziioannou *et al.* 2014; Klein *et al.* 2014), and for novel, embarrassingly parallel, parameter estimation methods (Pankow *et al.* 2015; Haster *et al.* 2015; Farr 2015). There is a critical need for low-latency parameter estimation methods which are able to account for all known features in the signals.

6. ACKNOWLEDGMENTS

We thank Atul Adhikari, Claudeson Azurin, Brian Klein, Brandon Miller, Leah Perri, Ben Sandeen, Jeremy Vollen, Michael Zevin for help running the MCMC analysis. T.B.L. and V.K. acknowledge NSF award PHY-1307020. B.F. was supported by the Enrico Fermi Institute at the University of Chicago as a McCormick Fellow. S.C. thanks the US-UK Fulbright Commission for personal financial support during this research period. Computational resources were provided by the Northwestern University Grail cluster through NSF MRI award PHY-1126812.

REFERENCES

- J. Abadie *et al.* (The LIGO Scientific Collaboration), (2014), arXiv:1411.4547 [gr-qc].
- F. Acernese *et al.* (Virgo Collaboration), Class.Quant.Grav. **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- H.-Y. Chen and D. E. Holz, Physical Review Letters **111**, 181101 (2013), arXiv:1206.0703.
- N. Cornish, L. Sampson, N. Yunes, and F. Pretorius, Phys. Rev. D **84**, 062003 (2011), arXiv:1105.2088 [gr-qc].

- T. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, *et al.*, Phys.Rev. **D85**, 082003 (2012), arXiv:1110.0530 [gr-qc].
- L. Sampson, N. Cornish, and N. Yunes, Phys.Rev. **D87**, 102001 (2013), arXiv:1303.1185 [gr-qc].
- E. E. Flanagan and T. Hinderer, Phys.Rev. **D77**, 021502 (2008), arXiv:0709.1915 [astro-ph].
- J. S. Read, C. Markakis, M. Shibata, K. Uryu, J. D. Creighton, *et al.*, Phys.Rev. **D79**, 124033 (2009), arXiv:0901.3258 [gr-qc].
- B. D. Lackey, K. Kyutoku, M. Shibata, P. R. Brady, and J. L. Friedman, Phys.Rev. **D89**, 043009 (2014), arXiv:1303.6298 [gr-qc].
- W. Del Pozzo, T. G. F. Li, M. Agathos, C. Van Den Broeck, and S. Vitale, Phys. Rev. Lett. **111**, 071101 (2013), arXiv:1307.8338 [gr-qc].
- L. Wade, J. D. E. Creighton, E. Ochsner, B. D. Lackey, B. F. Farr, T. B. Littenberg, and V. Raymond, Phys. Rev. D **89**, 103012 (2014), arXiv:1402.5156 [gr-qc].
- E. Berti *et al.*, (2015), arXiv:1501.07274 [gr-qc].
- K. Belczynski, T. Bulik, C. L. Fryer, A. Ruitter, J. S. Vink, *et al.*, Astrophys. J. **714**, 1217 (2010), arXiv:0904.2784 [astro-ph.SR].
- K. Belczynski, G. Wiktorowicz, C. Fryer, D. Holz, and V. Kalogera, Astrophys.J. **757**, 91 (2012), arXiv:1110.1635 [astro-ph.GA].
- I. Mandel, C.-J. Haster, M. Dominik, and K. Belczynski, (2015), 10.1093/mnras/slv054, arXiv:1503.03172 [astro-ph.HE].
- B. F. Schutz, Nature **323**, 310 (1986).
- D. E. Holz and S. A. Hughes, Astrophys. J. **629**, 15 (2005), arXiv:astro-ph/0504616 [astro-ph].
- S. Nissanke, D. E. Holz, S. A. Hughes, N. Dalal, and J. L. Sievers, Astrophys. J. **725**, 496 (2010), arXiv:0904.1017 [astro-ph.CO].
- D. Gamerman, *Markov Chain Monte Carlo: Stochastic Simulation of Bayesian Inference* (Chapman and Hall, London, 1997).
- J. Skilling, in *American Institute of Physics Conference Series*, American Institute of Physics Conference Series, Vol. 735, edited by R. Fischer, R. Preuss, and U. V. Toussaint (2004) pp. 395–405.
- F. Feroz and M. P. Hobson, (2007), arXiv:arXiv:0704.3704 [astro-ph].
- B. D. Metzger, G. Martínez-Pinedo, S. Darbha, E. Quataert, A. Arcones, D. Kasen, R. Thomas, P. Nugent, I. V. Panov, and N. T. Zinner, Mon. Not. Roy. Astron. Soc. **406**, 2650 (2010), arXiv:1001.5029 [astro-ph.HE].
- C. Pankow, P. Brady, E. Ochsner, and R. O’Shaughnessy, Phys. Rev. **D92**, 023002 (2015), arXiv:1502.04370 [gr-qc].
- C.-J. Haster, I. Mandel, and W. M. Farr, ArXiv e-prints (2015), arXiv:1502.05407 [astro-ph.IM].
- B. Farr, “Kombine,” (2015), <http://home.uchicago.edu/~farr/kombine/>.
- B. Farr, V. Kalogera, and E. Lijten, Phys. Rev. D **90**, 024014 (2014), arXiv:1309.7709 [astro-ph.IM].
- L. P. Singer and L. R. Price, (2015), arXiv:1508.03634 [gr-qc].
- L. P. Singer *et al.*, Astrophys. J. **795**, 105 (2014), arXiv:1404.5623 [astro-ph.HE].
- C. P. L. Berry *et al.*, Astrophys. J. **804**, 114 (2015), arXiv:1411.6934 [astro-ph.HE].
- B. Farr *et al.*, (2015), arXiv:1508.05336 [astro-ph.HE].
- T. Fragos and J. E. McClintock, ApJ **800**, 17 (2015), arXiv:1408.2661 [astro-ph.HE].
- B. Miller, R. O’Shaughnessy, T. B. Littenberg, and B. Farr, (2015), arXiv:1506.06032 [gr-qc].
- B. Farr, E. Ochsner, W. M. Farr, and R. O’Shaughnessy, Phys. Rev. **D90**, 024018 (2014), arXiv:1404.7070 [gr-qc].
- T. B. Littenberg, B. Farr, S. Coughlin, V. Kalogera, and D. E. Holz, Astrophys. J. **807**, L24 (2015), arXiv:1503.03179 [astro-ph.HE].
- J. Aasi *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), (2013), arXiv:1304.0670 [gr-qc].
- “LSC Algorithm Library suite,” <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.
- J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, *et al.*, Phys.Rev. **D91**, 042003 (2015), arXiv:1409.7215 [gr-qc].
- W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, *et al.*, Astrophys.J. **741**, 103 (2011), arXiv:1011.1459 [astro-ph.GA].
- F. Özel, D. Psaltis, R. Narayan, and J. E. McClintock, Astrophys.J. **725**, 1918 (2010), arXiv:1006.2834 [astro-ph.GA].
- L. Kreidberg, C. D. Bailyn, W. M. Farr, and V. Kalogera, Astrophys.J. **757**, 36 (2012), arXiv:1205.1805 [astro-ph.HE].
- S. Nissanke, M. Kasliwal, and A. Georgieva, Astrophys. J. **767**, 124 (2013), arXiv:1210.6362 [astro-ph.HE].
- C. Hanna, I. Mandel, and W. Voudsen, Astrophys. J. **784**, 8 (2014), arXiv:1312.2077 [astro-ph.HE].
- X. Fan, C. Messenger, and I. S. Heng, Astrophys. J. **795**, 43 (2014), arXiv:1406.1544 [astro-ph.HE].
- C. Cutler and E. E. Flanagan, Phys. Rev. D **49**, 2658 (1994), arXiv:gr-qc/9402014.
- C. L. Rodriguez, B. Farr, V. Raymond, W. M. Farr, T. B. Littenberg, *et al.*, Astrophys.J. **784**, 119 (2014), arXiv:1309.3273 [astro-ph.HE].
- S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, Phys. Rev. **X4**, 031006 (2014), arXiv:1308.3565 [gr-qc].
- A. Lundgren and R. O’Shaughnessy, Phys.Rev. **D89**, 044021 (2014), arXiv:1304.3332 [gr-qc].
- K. Chatzioannou, N. Cornish, A. Klein, and N. Yunes, Phys. Rev. **D89**, 104023 (2014), arXiv:1404.3180 [gr-qc].
- A. Klein, N. Cornish, and N. Yunes, Phys. Rev. **D90**, 124029 (2014), arXiv:1408.5158 [gr-qc].